



Table of Contents

- 2 Introduction
- 3 Memory Population Rules for Fully Balanced Configurations
- 6 CAPEX Saving from BigTwin's Design Advantage
- 7 Performance Improvements with BigTwin's Balanced Memory Design
- 9 Conclusions

White Paper

CAPEX Savings and Performance Advantages of BigTwin's Memory Optimization



Executive Summary

BigTwin is the 5th generation of Supermicro's award-winning Twin multi-node architecture design. It is the first multi-node dual-processor system in the world that features a no-compromise design to support 24 DIMM slots per node, typically only offered in 1U single-node server products.

Other high-performance multi-node systems on the market are usually designed with up to 16 DIMM slots per node, which prohibit a fully balanced memory configuration at the maximum memory capacity. Performance testing results have shown 52% improvements on BigTwin in a 16 DIMM configuration over other comparable systems, which could have significant impact on real-world applications.

Customers can also take advantage of the extra DIMM slots to achieve reduced CAPEX by acquiring less expensive DIMMs for the same total memory capacity. With 16 DIMMs, BigTwin offers the highest performance per dollar with its balanced memory configuration compared to other similar multi-node systems with unbalanced 16 DIMM configurations.

Super Micro Computer, Inc.
980 Rock Avenue
San Jose, CA 95131 USA
www.supermicro.com


Introduction

The Supermicro BigTwin platform offers the most powerful multi-node server solution that provides the most compute, memory, storage and power within a 2U rack space today. It is the 5th generation of Supermicro's award winning Twin multi-node architecture, and the first in the industry to offer 24 DIMM slots per node in a 2U 4-node design with support for free-air cooled data centers.

When compared with other high-performance multi-node systems that are designed with 16 DIMM slots per node, BigTwin's balanced memory architecture and higher maximum memory capacity allow customers to exploit significantly better memory throughput performance and achieve the highest performance/dollar benefits.

This paper shows how customers can achieve optimal memory performance and CAPEX in the following ways:

- Specific DIMM population rules shall be followed to ensure that a balanced memory configuration with dual memory channels on the BigTwin is applied to ensure optimal performance;
- The balanced memory architecture and additional DIMMs slots available on the BigTwin system allows it to outperform other similar systems for the same amount of memory installed, and at a lower price point when the price of DIMMs is factored in;
- Performance benchmarks have shown clear evidence of throughput improvements of a balanced memory system vs an unbalanced one.



1 of 4 dual-processor nodes



Figure 1. Supermicro 2U 4-Node BigTwin System (rear system view)

This paper is applicable to the 5th generation BigTwin servers supporting both the current and next generation Intel® Xeon® Scalable processors (codenamed Cascade Lake).

Memory Population Rules for Fully Balanced Configurations

Balanced memory configurations enable optimal interleaving across all attached memory channels so memory bandwidth can be maximized. The Intel® Xeon® Scalable platform is designed to provide up to 6 memory channels per CPU, and up to 2 DIMMs per channel. This is illustrated in the figure below:

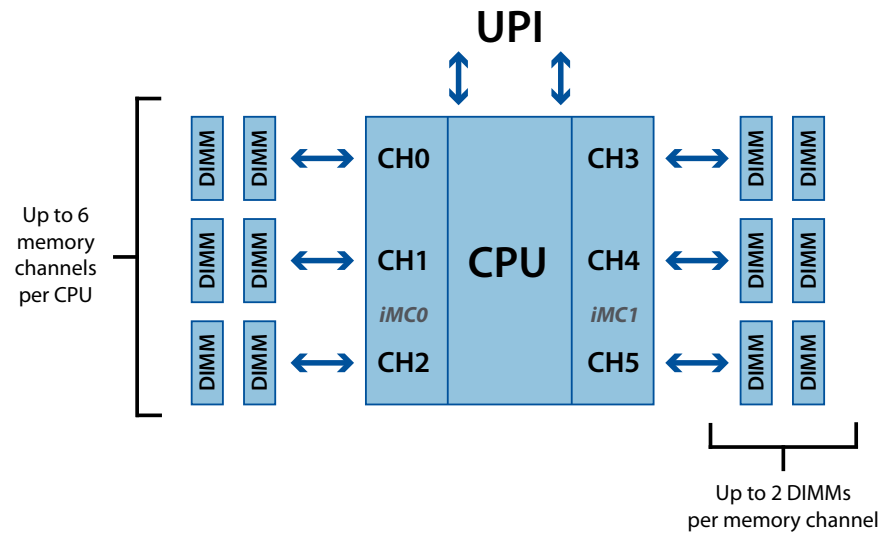


Figure 2. Memory Architecture of Intel Skylake-SP

Intel provided the following general guidelines prioritized in the following order:

1. Use identical DIMM types throughout the platform:
 - Same size, speed, and number of ranks
2. Maximize the same number of channels populated in each memory controller
3. Use a “balanced” platform configuration:
 - Populate memory channels equally
 - Identical DIMMs in all locations (size/speed/rank)
4. Use a “near-balanced” platform configuration:
 - Populate all memory controllers equally
 - Identical DIMMs in each “row”, but different sized DIMMs in row #1 vs. row #2

With this basic understanding, we can examine in more detail the architectural differences between BigTwin and other high-performance multi-node systems on the market today. Since the BigTwin can support full 6 memory channels with 2 DIMM slots per channel from each CPU, a fully balanced memory configuration can be achieved as below with identical DIMMs,




Figure 3. Example of BigTwin's Fully Balanced Memory Configuration When Populated with 12x 32GB DIMMs per CPU (per node)

BigTwin features 12 DIMM slots per socket that can maximize the potentials of the memory controller on the Intel® Xeon® Scalable platform. Other high-performance multi-node systems are usually designed with up to 8 DIMM slots per socket, which prohibit a fully balanced memory configuration with maximum capacity. This design compromise is illustrated in the diagram below,




Figure 4. A Fully Populated 8 DIMM slots per CPU Multi-Node System with an Unbalanced Memory Configuration (per node)

Since an unbalanced memory system can introduce a significant amount of performance degradation (see performance benchmarks section of this white paper), Supermicro has recommended DIMM configurations for each system or motherboard included in the user manuals to help customers optimize their system performance. The figure below shows the configurations recommended for the BigTwin.

Table 1. Possible DIMM Configurations on the BigTwin System (per CPU per node)

DIMM Population Recommendation for Both CPU Sockets on BigTwin												
	C1	C2	B1	B2	A1	A2	D2	D1	E2	E1	F2	F1
1 DIMM					●							
2 DIMMs					●			●				
3 DIMMs	●		●		●							
4 DIMMs			●		●			●		●		
5 DIMMs*	●		●		●			●		●		
6 DIMMs	●		●		●			●		●		●
7 DIMMs*			●	●	●	●		●		●		●
8 DIMMs			●	●	●	●	●	●	●	●		
9 DIMMs*	●	●	●	●	●	●		●		●		●
10 DIMMs*	●		●	●	●	●	●	●	●	●		●
11 DIMMs*	●	●	●	●	●	●	●	●	●	●		●
12 DIMMs	●	●	●	●	●	●	●	●	●	●	●	●

Supermicro Recommended DIMM Configurations for Optimal Bandwidth

All rows with an * illustrate imbalanced memory configurations.

CAPEX Saving from BigTwin's Design Advantage

Not only does BigTwin's unparalleled 24 DIMM slots per node design provide higher maximum memory capacity to enable a wider range of applications, customers can also take advantage of the dual channel memory slots to achieve reduced CAPEX by acquiring less expensive DIMMs for the same total memory capacity.

For example, based upon the cost of memory today, the sweet spot in terms of market price and capacity is 32GB DIMMs. To populate 512GB of total memory capacity per node, 16 DIMMs can be used.

Special Notes

- BigTwin has a 52% performance advantage over other systems with 16 DIMMs populated per node.
- 8 and 12 DIMM configurations per node are available through multiple vendors, but the total memory cost is higher because 64GB DDR4 2666 MHz memory modules are more expensive than 32GB DDR4 2666 MHz memory modules.

As shown in the previous section, memory performance is optimal with dual-channel balanced configurations. With BigTwin, a balanced configuration can either be 16 or 24 DIMMs installed per node.

With 16 DIMMs, BigTwin offers the highest performance per dollar with its balanced memory configuration compared to other similar multi-node systems with unbalanced 16 DIMM configurations.

When we look at the cost of memory today, 8x 64GB DIMMs is more expensive than 16x 32GB DIMMs. Consequently, the cost of 12x 64GB DIMMs is more expensive than 24x 32GB DIMMs. We believe the price delta and BigTwin's balanced DIMM configurations are the two main drivers for buying behaviors of these memory configurations per compute node. Of course, the total cost delta between 1x 64GB memory and 2x 32GB configurations become more evident at high server volumes, typically seen by our hyperscale data center customers.

Here's a simple equation to understand the cost delta between 64GB and 32GB DIMMs:

$$1x \text{ 64GB DIMM} = 2x \text{ 32GB DIMMs} + (\sim\$20 \text{ to } \sim\$30)$$

Therefore, the cost per node breaks down as:

$$8x \text{ 64GB DIMMs} = 16x \text{ 32GB DIMMs} + 8x (\sim\$20 \text{ to } \sim\$30) \text{ for } 512\text{GB per node}$$

$$12x \text{ 64GB DIMMs} = 24x \text{ 32GB DIMMs} + 12x (\sim\$20 \text{ to } \sim\$30) \text{ for } 768\text{GB per node}$$

Essentially, the total cost of delta between 64GB and 32GB DIMMs for a total of 512GB per node is about \$240 (8x \$30) at the upper limit. For 768GB per node, the cost delta per node is about \$360 (12x \$30). Furthermore, the total cost delta at the system level with 4 nodes becomes \$960 or \$1,440 respectively. For hyperscale deployments, the cost delta between memory types becomes significant.

Table 2. Summary of Competitive Solutions

	BigTwin (4-node)	Competitor A (4-node)	Competitor B (4-node)
Form Factor	2U / 4 Independent Nodes	2U / 4 Independent Nodes	2U / 4 Independent Nodes
Processors	Dual Intel® Xeon® Scalable Processors	Dual Intel® Xeon® Scalable Processors	Dual Intel® Xeon® Scalable Processors
Memory	24 DDR4-2666 DIMM slots, 3TB max per node	16 DDR4-2666 DIMM slots, 512GB max per node	16 DDR4-2666 DIMM slots, 1.5TB Max per node
Drive Bays	6x 2.5" with configurations for • 6 NVMe or • 4 NVMe/SAS/SATA + 2 SAS/SATA	6x 2.5" with configurations for • 2 NVMe/SAS/SATA + 4 SAS/SATA	6x 2.5" configurations for • 2 NVMe/SAS/SATA + 4 SAS/SATA
Internal Storage	• 2 M.2 NVMe/SATA slots • 1 SuperDOM port	• 1 M.2 SATA slot • 1 MicroSD slot	• 1 M.2 SATA slot • 1 MicroSD slot
PCI-E Expansion	• 2 PCI-E x16 slots • 1 SIOM (PCI-E x16) slot	• 1 PCI-E x16 slot • 1 OCP Mezz (PCI-E x16) • 1 Mezz for storage controller (PCI-E x8)	• 1 PCI-E x16 slot • 1 PCI-E x16 or FlexibleLOM (PCI-E x16)
Onboard NIC	SIOM	OCP Mezz	Dual 10GbE mLOM

Note: The cost of memory for either one of competitor's nodes with 768GB (12x 64GB) is ~\$355.20 more than BigTwin's nodes with 768GB (24x 32GB). For one complete system with 768GB of memory on each of the 4 nodes, customers can save ~\$1420.80 per system by selecting 32GB DIMMs instead of 64GB DIMMs.

Performance Improvements with BigTwin's Balanced Memory Design

The performance impact of balanced and unbalanced memory configurations were measured using the Stream Triad utility* with the following setup,

- Dual Intel® Xeon® Gold processor 6144
- 16x Samsung 32GB DDR4-2666 DIMMs
- Stream.icc17

* Stream Triad is the de facto industry standard benchmark for measuring memory bandwidth. Intel uses Stream-Triad to test performance at the system level and in the cloud to show the improved TCO and performance over previous CPU generations. (Source: <https://www.intel.com/content/www/us/en/benchmarks/xeon-scalable-benchmark.html>)

The test results are shown in the table below,

Table 3. Performance Comparison (16 DIMM Slots Populated)

Motherboards	24 DIMM-Slot Purely MB	16 DIMM-Slot Purely MB
Stream Triad Performance (MB/s)	142,890.70	75,371.90

- The 16 DIMM configuration showed a 52% performance improvement with the BigTwin's balanced memory design when compared to other multi-node systems with 16 DIMMs, which have unbalanced memory configurations.
- Intel also reported a ~50% gain in memory performance if motherboards have a balanced memory architecture, which supports 2+2+0, similar to the BigTwin's motherboard design.

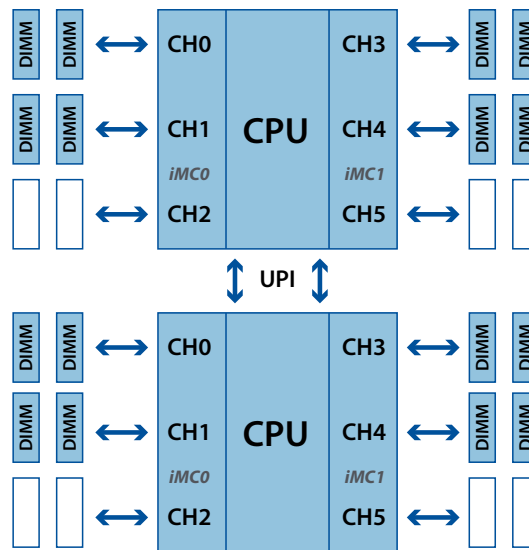


Figure 5. Another Example of BigTwin's Balanced Memory Configuration When Populated with 8 DIMMs per CPU (per node)

Conclusions

Supermicro BigTwin is the first multi-node platform introduced to the market featuring 24 DIMM slots per node that can fully exploit the potentials of the memory controller on the Intel® Xeon® Scalable platform. This unique design advantage is a result of engineering innovations at all system levels, including the compact Power Stick design.

Other high-performance multi-node systems are usually designed with up to 16 DIMM slots, which prohibit a fully balanced memory configuration with maximum capacity. Performance results have shown 52% performance improvements of BigTwin in a 16 DIMM configuration over similar systems which could have significant impact on real-world applications.

The additional DIMM slots allowing for balanced memory controllers on the BigTwin can potentially decrease datacenter CAPEX and provide extra flexibility for future upgrades. Since the real-world performance per dollar depends on how much memory is required per node, your Supermicro sales representative and solutions engineers can provide reference configurations at the sweet spot for your software stack.

For More Information

- Supermicro® BigTwin™ Solutions
<http://www.supermicro.com/bigtwin/>
- Supermicro® FatTwin™ Solutions
<http://www.supermicro.com/products/nfo/FatTwin.cfm>
- Supermicro® 1U TwinPro™
<http://www.supermicro.com/products/nfo/1UTwinPro.cfm>
- Supermicro® 2U TwinPro™ and 2U TwinPro2™
<https://www.supermicro.com/products/nfo/2UTwinPro.cfm>
- Supermicro 2U Twin2[®]
<http://www.supermicro.com/products/nfo/2UTwin2.cfm>
- Supermicro 2U Twin™
<http://www.supermicro.com/products/nfo/2UTwin.cfm>

About Super Micro Computer, Inc.

Supermicro® (NASDAQ: SMCI), the leading innovator in high-performance, high-efficiency server technology is a premier provider of advanced server Building Block Solutions® for Data Center, Cloud Computing, Enterprise IT, Hadoop/Big Data, HPC and Embedded Systems worldwide. Supermicro is committed to protecting the environment through its “We Keep IT Green™” initiative and provides customers with the most energy-efficient, environmentally-friendly solutions available on the market.

Learn more on www.supermicro.com

No part of this document covered by copyright may be reproduced in any form or by any means — graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system — without prior written permission of the copyright owner.

Supermicro, the Supermicro logo, Building Block Solutions, We Keep IT Green, SuperServer, Twin, BigTwin, TwinPro, TwinPro², SuperDoctor are trademarks and/or registered trademarks of Super Micro Computer, Inc.

Ultrabook, Celeron, Celeron Inside, Core Inside, Intel, Intel Logo, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Inside Logo, Intel vPro, Itanium, Itanium Inside, Pentium, Pentium Inside, vPro Inside, Xeon, Xeon Phi, and Xeon Inside are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

All other brands names and trademarks are the property of their respective owners.

© Copyright 2018 Super Micro Computer, Inc. All rights reserved.

